

Optimal Control Allocation with Load Sensor Feedback for Active Load Suppression, Experiment Development

Christopher J. Miller¹ and Dan Goodrick²
NASA Neil A. Armstrong Flight Research Center, Edwards, CA, 93523

The problem of control command and maneuver induced structural loads is an important aspect of any control system design. The aircraft structure and the control architecture must be designed to achieve desired piloted control responses while limiting the imparted structural loads. The classical approach is to utilize high structural margins, restrict control surface commands to a limited set of analyzed combinations, and train pilots to follow procedural maneuvering limitations. With recent advances in structural sensing and the continued desire to improve safety and vehicle fuel efficiency, it is both possible and desirable to develop control architectures that enable lighter vehicle weights while maintaining and improving protection against structural damage. An optimal control technique has been explored and shown to achieve desirable vehicle control performance while limiting sensed structural loads. The subject of this paper is the design of the optimal control architecture, and provides the reader with some techniques for tailoring the architecture, along with detailed simulation results.

Nomenclature

A	=	state derivative coefficients
Aero	=	aerodynamic
a_d	=	defined tracking constraint parameter
B	=	control effectiveness matrix
dap	=	roll stick
FAST	=	Full-scale Advanced Systems Testbed
FC	=	flight condition
fbk	=	feedback
filt	=	filtered
g	=	acceleration of gravity
H	=	trim tuning matrix
HM	=	hinge moment
I	=	inertia tensor
J	=	cost function value
K	=	roll mode gain
K_c	=	matrix of gains that result from the minimization of the cost function
KCAS	=	knots calibrated airspeed
k	=	iteration number
L	=	matrix of surface influence coefficients for each measured load
LOES	=	low order equivalent systems
la	=	left aileron
llef	=	left leading-edge flap
lr	=	left rudder
ls	=	left stabilator
ltef	=	left trailing-edge flap

¹ Full-Scale Advanced Systems Testbed Chief Engineer, Flight Controls and Dynamics, P.O. Box 273 Edwards, California/MS 4840D, AIAA Member.

² Research Engineer, Instrumentation Branch, P.O. Box 273 Edwards, California/MS 4800-1409, AIAA Member.

M	=	measured loads vector
m	=	mass
N	=	load cost function simplification parameter
NDI	=	nonlinear dynamic inversion
n	=	load constraint steepness
OCLA	=	optimal control and load allocation
Per	=	persistence counter
PI	=	proportional plus integral
p	=	roll rate
psf	=	pounds per square foot
q	=	pitch rate
q _c	=	impact pressure
R	=	relaxation factor
r	=	yaw rate
ra	=	right aileron
ref	=	reference
rlef	=	right leading-edge flap
rr	=	right rudder
rs	=	right stabilator
rtef	=	right trailing-edge flap
S	=	aerodynamic reference area
SMI	=	structural modal interaction
s	=	Laplace operator
TEF	=	trailing-edge flaps
u	=	control surface vector (m-measured, p-trim)
x	=	state vector
α	=	angle of attack
β	=	angle of side slip
γ	=	load constraint weight
ε	=	trim cost weight
ω	=	angular velocity vector
ω_n	=	roll mode natural frequency
ζ	=	damping ratio

Superscripts

\rightarrow	=	vector quantity
\cdot	=	time derivative
b	=	body fixed axis
e	=	earth fixed axis

I. Introduction

CURRENT control applications account for structural load limits by: limiting the types of control algorithms that can be applied to a given application, requiring high structural margins resulting in less efficient designs, and placing procedurally enforced restrictions on pilot control actions and maneuvers. One technology that may alleviate some of these constraints is to utilize structural sensing feedback to actively sense and limit critical loads. A number of theoretical and simulation studies¹⁻³ have been published on this topic, but a flight experiment is essential to evaluate the feasibility and applicability of the approach. This technology would facilitate lighter weight optimized structure, increase robustness of the system to off nominal and unpredicted situations (stall/spin, accidental damage), and pave the way for highly optimized high performance control algorithms with less intuitive structural interaction.

The specific goal for the experiment that is the subject of this paper is to explore the merits of optimal control allocation³ with structural feedback in flight. Flight on a full-scale piloted vehicle will provide valuable insight into the performance of the technique in a flight environment and help explore the interactions between the pilot and the approach. This paper deals with the experiment formulation and the nonlinear simulation predictions.

II. Problem Statement and Research Objectives

The current design paradigm for flight controls and vehicle structures has worked extremely well, and resulted in exceptional levels of safety and performance. However, as American Airlines Flight 587 (Fig. 1) illustrates, there are limits to the existing design approach.⁴ Flight 587 encountered the wake of a larger transport, the result of which violently upset the airplane. The pilot, acting consistent with his vehicle upset training, attempted to recover control with the aggressive use of the rudder pedals. The combination of maneuver and rudder induced loads broke the vertical tail off, resulting in complete loss of control. All of this activity occurred below the published maneuvering speed of the airplane which had been incorrectly understood as a safe speed for rapid and aggressive rudder pedal inputs. This incident provides a sobering illustration of the possible benefits of a control architecture that protects aircraft structure while using all available control effectors to provide the pilot with the necessary control response. Future aircraft with load limiting control technologies could give the pilot the best possible control response in an emergency while minimizing the threats posed by control and maneuver induced catastrophic structural failure.



Figure 1. American Airlines Flight 587.

The Optimal Control and Load Allocation (OCLA) experiment is designed to test the effectiveness and robustness of an optimal control allocator with strain-gauge feedback at limiting aircraft loads while producing the desired aircraft response commanded by the pilot. Specifically, the experiment is designed to limit the measured aileron hinge moment on the National Aeronautics and Space Administration (NASA) Full-Scale Advanced Systems Testbed (FAST) while maintaining desired handling qualities for an air-to-air tracking task. The following is a list of the detailed experiment objectives.

Objective 1: Limit the aileron motion subject to a defined load constraint.

Metric: The measured aileron hinge moment does not exceed 100% of the specified value for any of the test scenarios.

Objective 2: Maintain the roll performance of the original controller that does not utilize structural load as a constraint.

Metric: No worse than a 10% reduction in the roll mode time constant or the steady state roll rate for up to half roll stick maneuvers.

Objective 3: Maintain the handling qualities ratings of the original controller that does not utilize structural load as a constraint.

Metric: No worse than one Cooper-Harper⁵ rating degradation for the tasks with aileron motion clipping due to the load constraint.

III. Scope and Technical Approach

The NASA Armstrong FAST was used as an in-flight laboratory to demonstrate the technical feasibility of load limiting control via optimal control allocation. The testbed aircraft and the research envelope have been specifically designed to facilitate rapid prototyping of advanced control and sensor technologies. Within the limited research envelope (Fig. 2) a wide array of sensor technologies and control schemes can be evaluated in flight with minimal testing between design iterations.⁶ This envelope is commonly referred to in other FAST related publications as the Class B envelope.




Figure 2. Research envelope (Class B).

The OCLA experiment used the outputs from four foil strain bridges as feedbacks, two redundant bridges on each aileron. The redundant sensors were compared to the primary strain gauges to mitigate concerns about failures of research instrumentation used as feedback in a control law. When the strain gauges were installed they were not intended as primary flight control feedback parameters and as such are not as reliable as would be desired for a production representative design. These measured strains on the aileron control arms were used to calculate the actual hinge moments on each aileron. Excessive hinge moments can cause premature fatigue and failure of the ailerons and the surrounding structure. Conceptually, the experiment is designed to illustrate that the optimal allocator will actively limit the measured aileron hinge moments and reallocate the pilot commanded body rates to other control surfaces such as the trailing-edge flaps and the stabilators, thus, sensing and limiting the imparted aileron loads, but still giving the pilot the commanded vehicle response.

Due to the low dynamic pressures achievable in the Class B envelope, it is not possible to achieve aileron hinge moments approaching the actual limit for the ailerons; therefore, it is necessary to set an arbitrarily low target limit to test out the algorithm. This reduced limit was tunable in flight and was based on the largest hinge moment measured during past experiments within the Class B envelope of 15000 in-lb.

IV. Detailed Design Description

The OCLA experiment is built on top of the nonlinear dynamic inversion (NDI)^{7,8} baseline controller utilized by previous experiments flown on the FAST aircraft. The NDI is based on a conservation of angular momentum formulation where pilot commands are passed through reference models with predicted desirable handling qualities. Those reference command outputs are used as the commands to the allocation scheme that utilizes aerodynamic table lookups to determine the control moment capabilities for each of the available control surfaces, and allocates the angular acceleration commands to those surfaces based on the desired allocation scheme. Additional components of

the control law that are important to understand, but not the subject of this paper are: the proportional plus integral (PI) compensator which is used to improve the robustness of the closed loop design, and the structural notch filters on the body rate feedbacks (Fig. 3). It is worth noting here that some of the experiment configurations did not contain any structural notch filtering on the strain feedbacks. This omission was deliberate to see what kinds of structural modal interactions (SMI) would be excited with the OCLA approach. However, there were also available configurations that did contain minimal filtering of the strain feedbacks to mitigate possible adverse control-structural interactions and reduce the probability of a loss of mission due to undesirable SMI.

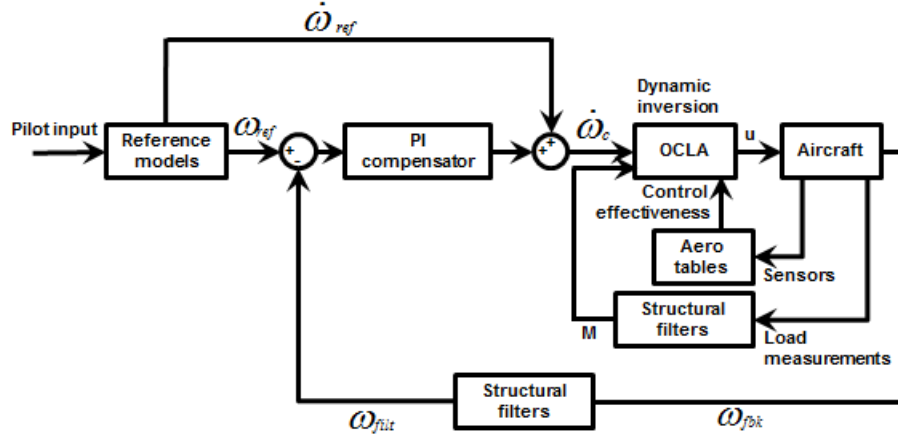


Figure 3. Experiment block diagrams.

A. Cost Function for an Optimal Control Allocator with Load Constraints and Tuning Guidance

The derivation of the optimal control allocator that is the subject of this paper is based on a dynamic inversion framework which starts with conservation of angular momentum [Eq. (1)], where the commands to the control allocator are angular accelerations.

$$\dot{\overline{Hm}}^e = \sum \overline{Mm} \quad (1)$$

Equations (2) through (9) show the derivation of the tracking performance cost function element from the conservation of angular momentum from Eq. (1). This element of the cost function will be discussed in relation to the other elements later in this section.

$$\overline{Hm} = I\vec{\omega} \quad (2)$$

$$\dot{\overline{Hm}}^e = I\dot{\vec{\omega}}^b + \vec{\omega}^b \times I\vec{\omega}^b \quad (3)$$

$$\sum \overline{Mm} = q_c S A x + q_c S B u \quad (4)$$

$$x = \begin{bmatrix} 1 \\ \alpha \\ \beta \\ p \\ q \\ r \end{bmatrix}, u = \begin{bmatrix} la \\ ra \\ ls \\ rs \\ lr \\ rr \\ llef \\ rlef \\ ltef \\ rtef \end{bmatrix} \quad (5)$$

$$I\dot{\vec{\omega}}^b + \vec{\omega}^b \times I\vec{\omega}^b = q_c S A x + q_c S B u \quad (6)$$

$$B u = \frac{1}{q_c S} [I\dot{\vec{\omega}}^b + \vec{\omega}^b \times I\vec{\omega}^b] - A x \quad (7)$$

$$a_d \stackrel{\text{def}}{=} \frac{1}{q_c S} [I\dot{\vec{\omega}}^b + \vec{\omega}^b \times I\vec{\omega}^b] - A x \quad (8)$$

$$\therefore \text{desire } \min(\|B u - a_d\|_2^2) \quad (9)$$

Desired tracking performance can be achieved by utilizing a simple and straight forward weighted pseudo inverse of B, as was done in Refs. 7 and 8. However, the weighted pseudo inverse approach does not account for surface rate or position limits, and it does not allow for the addition of constraints such as load or wing shape. Therefore, it becomes desirable to define an alternate cost function, one that yields formulations that permit the inclusion of surface limits, load constraints, and other optimization constraints in addition to command tracking. The cost function described in Fig. 4 was developed as part of this research effort to meet these objectives.




Figure 4. Cost function overview.

Where the first L2 norm term in the cost function minimizes tracking error (J tracking), the second drives the surfaces to their desired trim positions (J trim) (commonly in the null space of B), and the last one represents the minimization of a load constraint designed to provide a hard constraint at 100% of the design load (J load). The L2 norm was chosen for use in the cost function for its computational simplicity, and because it is easily twice differentiable. The L2 norm also yields the desirable feature that its second derivative with respect to the command

vector (u) is positive definite as long as the combination of L , H , and B is full rank. Therefore, this cost function formulation forms a convex space with a global minimum. Other norms could be used, but would require different optimization techniques to solve for their minima that may or may not have the property of being a convex space. Figure 4 provides a complete overview of the design objective being satisfied by each component within the cost function.

1. *Command Tracking Cost Function Term (J tracking)*

The primary objective for any control allocator is that it must provide surface commands that track the desired input commands. The first L2 norm within the cost function shown in Fig. 4 represents this command tracking objective. Under nominal operational scenarios when the pilot is performing aircraft maneuvers, this element should be the dominant objective in the cost function. Low values of this aspect of the cost function are required to achieve desired handling qualities characteristics. The actual formulation of this aspect of the cost function comes directly from the conservation of angular momentum formulation of dynamic inversion discussed at the beginning of this section.

2. *Trim Constraint Cost Function Term (J trim)*

It is important that any optimal control allocator provide a means for a designer to specify a desired solution within the over determined space spanned by B for the steady state trim positions for each surface. The trim L2 norm term is analogous to a leakage term that should only be the dominant objective when all the other terms are small. Keeping the weight (ϵ) on J trim as small as practical is desirable and improves the command tracking and load limiting performance; however, without this term the surfaces trim in odd and not intuitively obvious ways.

The trim surface commands (u_p) can be anything in the null space of the control effectiveness matrix (B). It is important to specify a trim command to give the vehicle desirable control characteristics, to put the surfaces in positions where they will have sufficient control power to generate the necessary control moments, and to reduce aerodynamic drag. A trim optimizer has been developed that reduces fuel flow and has been flown on the FAST vehicle which could be used to determine desirable trim solutions.^{9,10} However, the OCLA experiment was designed to mimic the production trim scheme which uses trailing-edge flaps and leading-edge flaps to reduce the angle of attack of the aircraft for elevated g maneuvers.

The scalar weight (ϵ) is used for tuning how aggressively the cost function attracts the control surfaces back to their specified trim positions. It must be small enough to allow the surfaces to move enough to track the desired dynamics, but large enough to attract the surfaces back to their desired trim positions in the absence of large commands. Tuning for this scalar value in the experiment for this paper was done in the nonlinear simulation by trial and error. A range of orders of magnitude for this scalar were tested in flight to see how changes in this parameter affect load limiting and handling qualities.¹¹

The square matrix H was used to tune the relative importance of maintaining the commanded trim position for each surface individually. For instance for the FAST, it is much more important for the trailing-edge flaps (TEF) to be close to their trim command than it is for the ailerons. This behavior is desirable due to the fact that the TEF are slow, and maintaining the TEF close to their commanded trim position is important for keeping angle of attack low, while the ailerons are fast and primarily used for tracking roll commands. The matrix H allows the designer fine control of the surface usage priorities without dedicating any of the surfaces just to the trim, command tracking, or load limiting roles. Tuning the values in this matrix is straight forward. For simplicity, the H matrix for this experiment was tuned empirically in the nonlinear simulation to yield surface usages that resemble the production control allocation for the host vehicle.

3. *Load Constraint Cost Function Term (J load)*

The primary design goal for the cost function load constraint for this experiment was to provide a practical hard constraint that would prevent any of the measured loads (M) from exceeding 100% of their specified limits. A very important secondary design goal was that the cost function allow free use of the control surfaces below a specified percentage of the maximum allowable load. Specifically for this implementation, the goal was to keep the aileron hinge moments below 100% of the specified limit 100% of the time, but to allow free use of the ailerons as long as their individual loads were below 80% of the specified limit.

The load constraint in the cost function assumes a simple relationship between incremental load and incremental surface deflection. This assumption is commonly made even for fundamentally nonlinear relationships because it is assumed that there exists a small region within which any relationship can be assumed linear. The region where this assumption is valid may be small for some applications, but was found to work exceptionally well for the mostly linear

relationship between aileron hinge moment and aileron position. M is a vector of the measured loads, and L is the matrix of incremental load coefficients mapping the influence of each control surface to each load. The L matrix could be a table look-up based on the current aircraft states and surface positions similar to the A and B matrices for the tracking constraint; however, a constant matrix scaled by impact pressure was found to be sufficiently accurate for the purposes of this experiment. It is also important to normalize each load and the associated surface influence coefficients by their associated limit load. This normalization prevents the largest load from dominating and allows the loads that are the closest to their associated limits to dominate the load cost constraint. Further shaping of individual loads may be desirable. In Ref. 11, it was discovered that for the current formulation when multiple loads get near 80% of their limits that the combination of loads has an adverse effect on handling qualities.

γ is a scalar weight used for setting the relative importance of the load constraint with respect to the other control objectives. γ must be tuned in conjunction with the load constraint steepness (n) to tailor at what load level this constraint dominates the control response. The load constraint steepness (n) is tuned to provide a steeper load constraint at higher load and a nice flat nearly zero one at lower load. These two parameters allow this constraint to act like a hard load constraint at 100% of the limit load while giving it the mathematical simplicity of a continuous load minimization constraint. Tuned together they provide the designer with the flexibility to tune the cost function such that when loads are low, the constraint is functionally dormant, but when the loads are high it is the dominant constraint. Higher values of n can exhibit convergence issues with a Newton-Rapson optimization approach (discussed in section B) because it introduces nonlinearity into the shape of the cost function, but if the value of n is too low the constraint behaves less like a hard constraint and is more active at lower loads. This formulation provides a practical hard constraint on the load without requiring a unique mapping from loads to surface positions.

A useful way to look at how γ and n can be tuned is to consider the loop gain associated with surface usage versus measured load. For the cost function above, the gain associated with every input to surface command can be determined by the following equation, derived from $\frac{\partial J}{\partial u} = 0$. The derivation of which and justification for that being the global minimum are discussed in section B. Equation (10) represents the u that solves $\frac{\partial J}{\partial u} = 0$ for fixed values of N , M , and a_d .

$$u = K_c^{-1} B^T a_d - K_c^{-1} \gamma n N^{n-1} L^T M \quad (10)$$

Equations (11) and (12) expand out the details of K_c and N in Eq. (10).

$$K_c = B^T B + \varepsilon H^T H + \gamma n N^{n-1} L^T L \quad (11)$$

$$N = \|M + L(u - u_m)\|_2^2 \quad (12)$$

The analysis and plots that follow show the effects of the various tuning features of the cost function on the individual values of K_c for varying values of N . The effect of tuning γ can be seen in Fig. 5 and Fig. 6. Higher values of γ result in more aggressive load limiting (higher gain in Fig. 5). However, higher values of γ restrict aileron usage at lower loads (fading out of ailerons earlier in Fig. 6). Stated simply, it is desirable to have high γ from a load limiting perspective, but undesirable because it results in aileron limiting at too low of a measured load and adversely affects handling qualities.




Figure 5. Load constraint weight comparison for $n=4$ for gain on load feedback to aileron command versus resulting hinge moment.

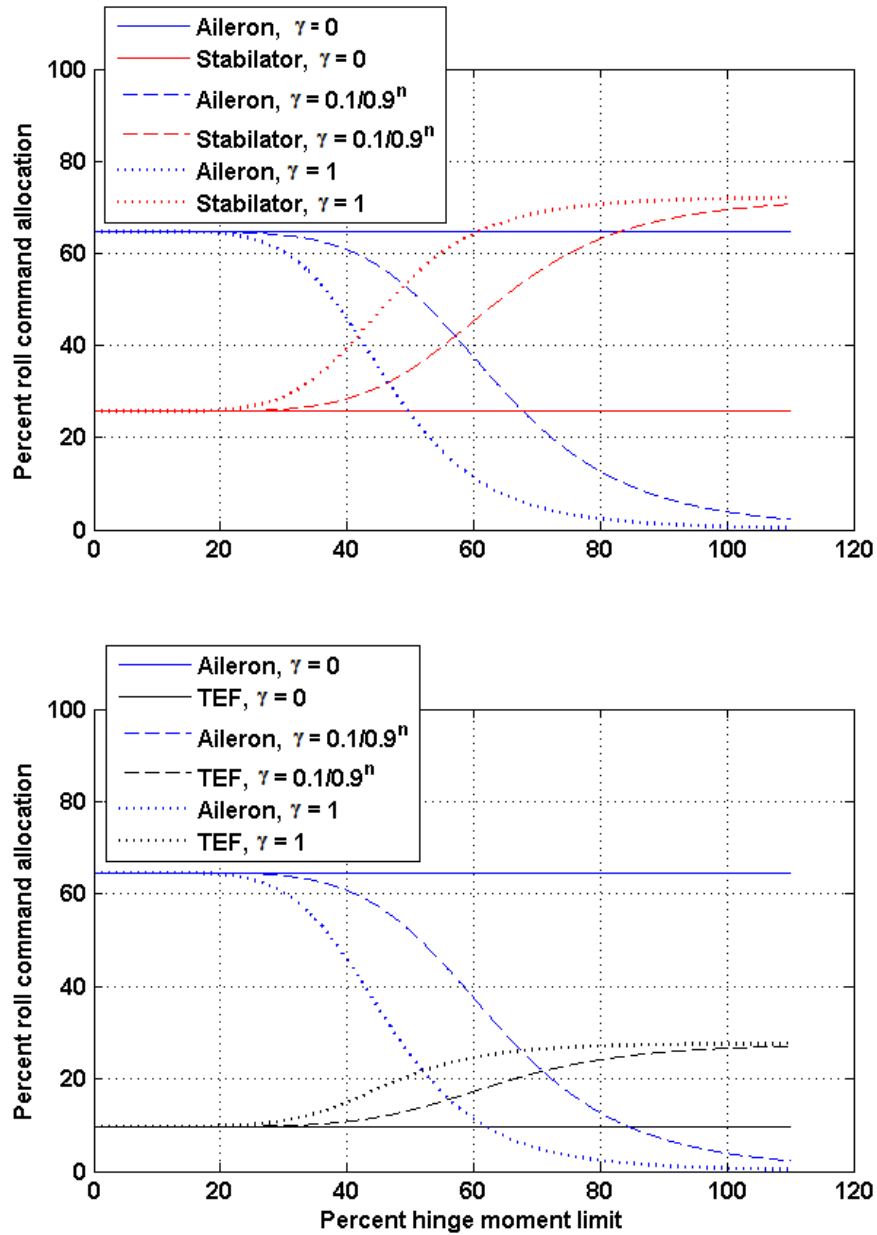


Figure 6. Load constraint weight comparison for $n=4$ of control surface allocation for roll command versus resulting aileron hinge moment.

Figures 7 and 8 illustrate the effect of the load constraint steepness (n) on aileron usage and control allocation. Higher n values exhibit more aggressive load limiting at higher load (higher gain in Fig. 7) and less aggressive load limiting at lower load (lower gain in Fig. 7). Additionally, steeper load constraints (higher n) result in delaying the transition from ailerons used for roll control to ailerons used for load limiting (Fig. 8). Both of these behaviors are highly desirable and help meet the objectives of limiting the measured load to below 100% while allowing free use of the ailerons at lower load. The steep load constraints make it possible to achieve good handling qualities while still using the ailerons to actively limit the measured hinge moment.




Figure 7. Load constraint steepness comparison for gain on load feedback to aileron command versus resulting hinge moment ($\gamma=0.1/0.9^n$).




Figure 8. Load constraint steepness comparison of control surface allocation for roll command versus resulting aileron hinge moment ($\gamma=0.1/0.9^n$).

B. Cost Function Optimization Technique

Optimization of the cost function outlined in section A can be accomplished in any number of ways. The formulation of the cost function is such that $\frac{\partial^2 J}{\partial u^2}$ is positive definite as long as the B, H, and L matrices are formulated correctly, which results in the solution for $\frac{\partial J}{\partial u} = 0$ being a global minimum of the cost function. Newton-Rapson [Eq. (13)] was the approach chosen for solving $\frac{\partial J}{\partial u} = 0$ for this experiment due to its simplicity and the fact that it is computationally efficient and was found to have desirable convergence properties for the shape of the cost function.

$$u_{k+1} = u_k - R * \frac{\partial^2 J^\dagger}{\partial u^2} \frac{\partial J}{\partial u} \quad (13)$$

The following is the derivation of the first and second derivatives of the cost function [Eq. (14)] for use within the Newton-Rapson iterative optimization scheme.

$$J = \|Bu - a_d\|_2^2 + \varepsilon \|H(u - u_p)\|_2^2 + \gamma [\|M + L(u - u_m)\|_2^2]^n \quad (14)$$

Expanded out, Eq. (14) turns into Eq. (15).

$$J = (Bu - a_d)^T (Bu - a_d) + \varepsilon (u - u_p)^T H^T H (u - u_p) + \gamma [(M + L(u - u_m))^T (M + L(u - u_m))]^n \quad (15)$$

Equation (16) is the result of multiplying out all of the elements of Eq. (15) and collecting like terms.

$$J = [u^T B^T Bu - u^T B^T a_d - a_d^T Bu + a_d^T a_d] + \varepsilon [u^T H^T Hu - u^T H^T H u_p - u_p^T H^T Hu + u_p^T H^T H u_p] + \gamma [M^T M + M^T L(u - u_m) + (u - u_m)^T L^T M + (u - u_m)^T L^T L(u - u_m)]^n \quad (16)$$

The first derivative of Eq. (16) can be seen in Eq. (17).

$$\frac{\partial J}{\partial u} = [B^T Bu + (B^T B)^T u - B^T a_d - (a_d^T B)^T] + \varepsilon [H^T Hu + (H^T H)^T u - H^T H u_p - (u_p^T H^T H)^T] + \gamma n [M^T M + M^T L(u - u_m) + (u - u_m)^T L^T M + (u - u_m)^T L^T L(u - u_m)]^{n-1} [(L^T M)^T + L^T M + L^T L(u - u_m) + (L^T L)^T (u - u_m)] \quad (17)$$

Equation (18) is the result of collection all of the like terms in Eq. (17).

$$\frac{\partial J}{\partial u} = 2\{B^T [Bu - a_d] + \varepsilon H^T H (u - u_p) + \gamma n N^{n-1} L^T [M + L(u - u_m)]\} \quad (18)$$

Taking the derivative of Eq. (18) results in Eq. (19) which is the final element of Eq. (13) used to minimize the cost function J in Eq. (14). Equation (20) is an expansion of the term N.

$$\frac{\partial^2 J}{\partial u^2} = 2\{B^T B + \varepsilon H^T H + 2\gamma n(n-1)N^{n-2}[L^T M + L^T L(u - u_m)][L^T M + L^T L(u - u_m)]^T + \gamma n N^{n-1} L^T L\} \quad (19)$$

$$N = M^T M + M^T L(u - u_m) + (u - u_m)^T L^T M + (u - u_m)^T L^T L(u - u_m) = \|M + L(u - u_m)\|_2^2 \quad (20)$$

Convergence properties are important for any iterative scheme, especially for a flight critical application. Under most operating conditions the Newton-Rapson approach as applied converged in three or less iterations; however, achieving good convergence with large values of the load constraint exponent for aggressive maneuvers proved a bit more challenging. A self-adjusting relaxation factor (R) was found to address the poor convergence issues that arose due to the steepness of cost function at high load with large n values. Because the cost function forms a convex space, the value of the cost function should monotonically decrease with each iteration of the optimization routine. If the value of the cost function increases for an iteration, the step size for that iteration was too large. Therefore if J increases

for an iteration, the relaxation factor (R) is reduced so that a smaller step is taken and the iteration is repeated. Each iteration is repeated with smaller and smaller relaxation factors (smaller steps) until J is reduced by the iteration.

A number of other algorithm features were necessary to account for the rigorous demands placed on the software because it was flight critical. The rank of the control effectiveness matrix (B) was checked using a minimum singular value method to assure that the cost function was well conditioned for each iteration. The condition number of the second derivative of the cost function with respect to u, which is a square matrix, was evaluated to prevent inversion of a poorly conditioned matrix. Optimization convergence was verified by both the value of the cost function and the norm of the first derivative of the cost function; both needed to be below a specified value for the algorithm to be considered adequately converged. The number of iterations allowed to reach convergence was limited to avoid over running the allotted computation time within the flight hardware for cases where convergence may fail or take excessive iterations. The research control law was disengaged and control was handed back to the production control law if any of these checks fails. Further research is necessary to determine how to implement this control approach on a platform without a robust production controller to fall back on.

Surface rate and position limits can be enforced between iterations (rate limits can be implemented as floating position limits for each individual frame). However, implementing rate limits in this way requires flight critical instrumentation of control surface positions with minimal delay and sufficiently high sample rates. The fundamental principal was verified in the simulation and shows promise for future implementations. However, due to limitations in the production measurements for control surface positions, surface rate limiting was not implemented for this experiment. The production measurements for aileron position were used within the load constraint, and they were found to have adequate sample rate and delay properties.

V. Hinge-Moment Sensor Model and Signal Management

This experiment uses the outputs from four foil strain bridges measuring aileron hinge moments. Sensor characterization and models of these sensors were needed for use in the simulation for design and testing the algorithms prior to flight. Models were developed from flight data from previous research flights within the same research envelope planned for this experiment. The models were developed using data from a variety of maneuvers including: stick and rudder doublets, half stick 360° rolls, 2-g loaded rolls (rolling pullout), and 2.5-g level turns. A linear least squares approach was used to determine a model that was a best fit to the data [Eq. (21)].

$$hm = q(C_\alpha \alpha + C_{ail} ail + C_p \frac{pb}{2V} + C_0) \quad (21)$$

While evaluating the quality of the resulting fit, there appeared to be an additional dependence on angle of attack. An additional third order polynomial of hinge moment from angle of attack was identified from the residual error which improved the overall fit [Eq. (22)].

$$hm = q \left(C_\alpha \alpha + C_{ail} ail + C_p \frac{pb}{2V} + C_0 \right) + P_3 \alpha^3 + P_2 \alpha^2 + P_1 \alpha + P_0 \quad (22)$$

Figure 9 shows an example of how well the above model characterizes the sensed hinge moment for a number of different maneuvers. However, Fig. 9 also shows that there is a substantial increase in the high frequency content of the sensed hinge moment for both the 2.5-g level turn and the loaded roll. A detailed investigation revealed that the high frequency peaks in the data for both hinge-moment sensors on a given wing were time correlated. This correlation means that the high frequency content represents actual load related to aerodynamic buffet and not just sensor noise. The fact that the content represents real load means that the high frequency content must be accounted for and cannot just be filtered out of the feedback signal. Modeling this behavior in the simulation was accomplished by adding a Gaussian random noise parameter with a standard deviation that depends on angle of attack to the hinge moment from the model above. Figure 10 shows the model fit for the 2.5-g level turn with the buffet model implemented. Figure 10 shows that the model reasonably approximates the aerodynamic buffet effect on the hinge-moment measurement. This model was used to determine the predicted performance of the optimal allocation scheme in the simulation and to tune all of the cost function and optimization related parameters.

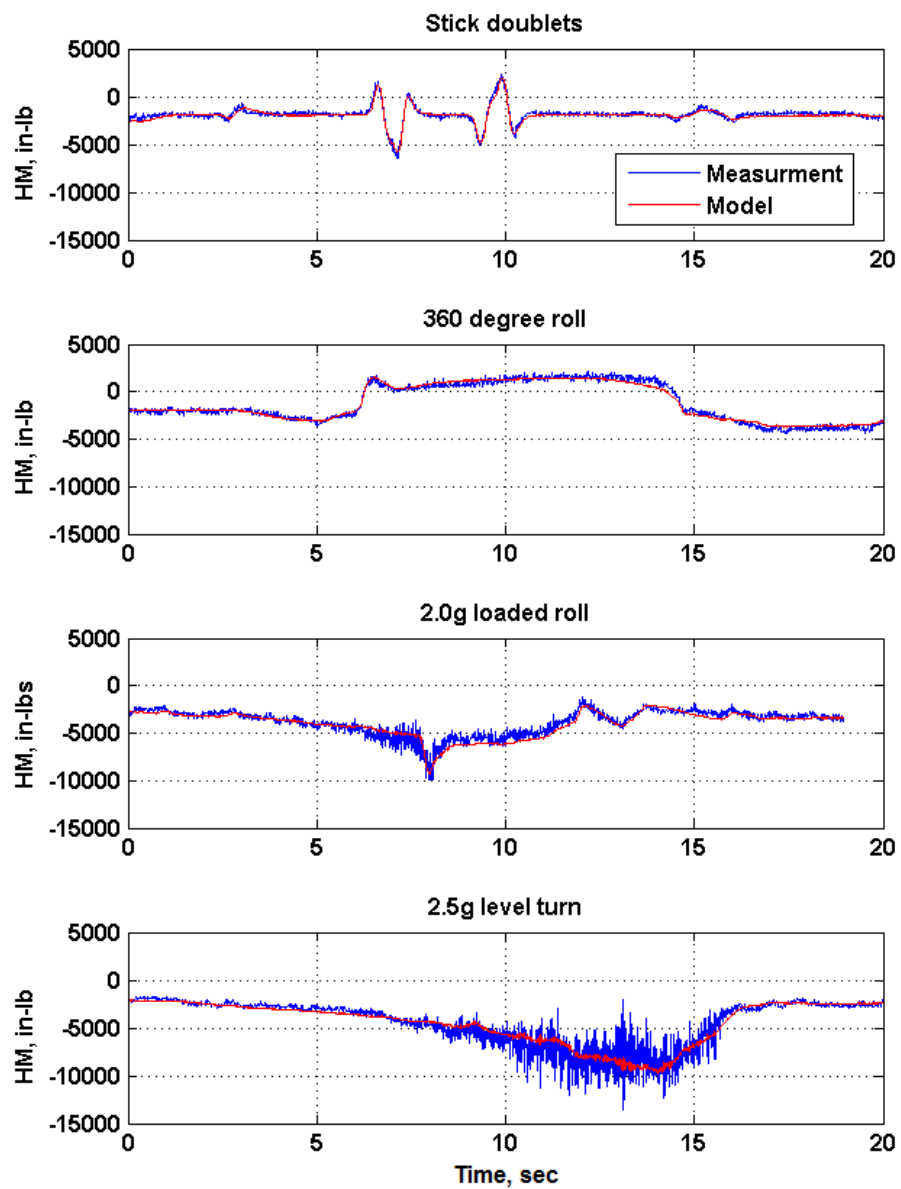


Figure 9. Aileron hinge-moment model illustration (no buffet model).




Figure 10. Aileron hinge-moment model illustration for elevated g with buffet model.

For OCLA the strain sensors on both ailerons are used as feedbacks to the control system; therefore, integrity validation is a crucial component of the signal processing design. These sensors were not originally intended for use as control feedbacks and as such were not installed with that level of robustness in mind. For the purposes of this research, which is to explore the feasibility of the control architecture, the sensors have sufficient reliability for a short flight-test campaign. However, it was still necessary to implement validity checks to verify the integrity of the sensed strains. Fiber optic strain sensing or some other advanced load sensing technique may be better suited for feedback control for a production vehicle, and may be a topic for future research.

Each aileron has two measurements of the hinge moment, one on the inboard hinge and one on the outboard hinge. These two measurements can be considered redundant because the aileron control surface can be assumed to be a rigid surface for the dynamic pressures and Mach numbers within the class B envelope. The redundant measurements are compared with each other, and if they do not agree within some tolerance, then at least one of the sensors has failed. The flow chart in Fig. 11 shows the detailed decision logic used to assess if a sensor has failed, determine which one has failed, and then output a voted hinge-moment value for use within the controller. The voting scheme uses red and yellow limits with a persistence counter to check the sensors against each other. If the sensors agree with each other they are averaged and that average is used as the control feedback. If they do not agree it is assumed that at least one of the sensors is failed. The values are then checked against their respective models, and the one closest to its respective model (without the buffet model) is used as the control feedback. An additional mitigation to prevent closed loop control on a failed sensor is to utilize a valid range check before using the sensor for control. If the voted value of the sensor is outside of a valid range, both sensors are assumed failed and the test is aborted, the system disengaged, and control is returned to the production controller. This voting scheme helped to maximize mission success. In the event that at most one sensor on each aileron failed, the system would declare it failed and use the other good sensor, approximating a fail-operational redundant system.

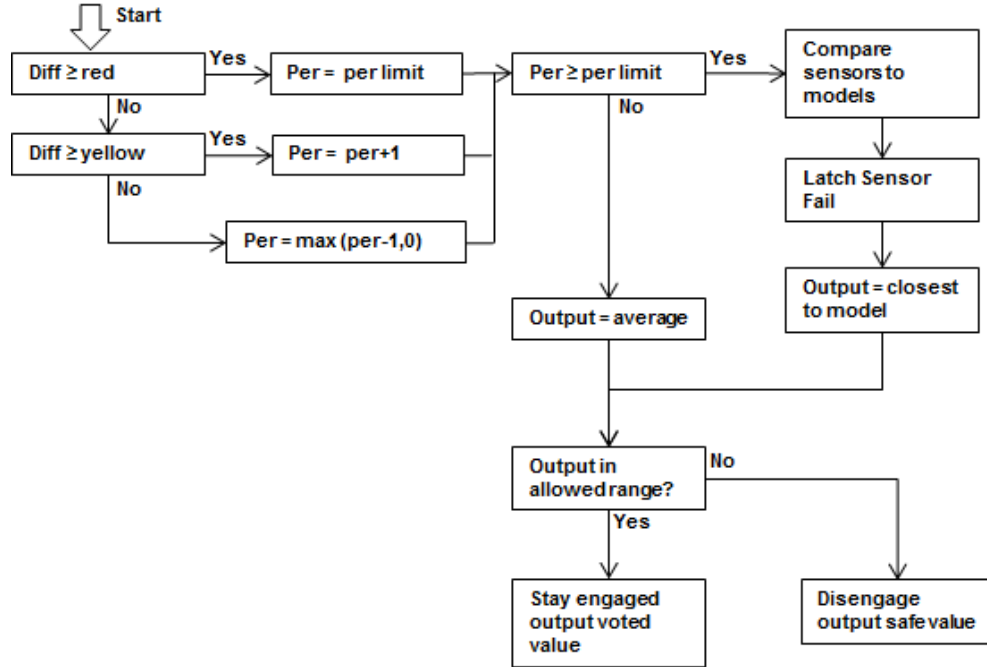


Figure 11. Redundancy management flow chart.

Accounting for vehicle structural dynamics and their interactions with the control laws is an important aspect of any real world control system design. The strain gauges in the optimal control allocation scheme are feedback sensors; therefore, the loops that contain them must account for structural dynamics. However, due to the unconventional feedback approach, it is not clear that the tried and true approaches for SMI analysis will be applicable. The optimization of a cost function does not lend itself nicely to the traditional loop gain and phase analysis for determining the stability properties of the strain feedback loops. Additionally, models for the dynamic structural modes for the aileron hinge-moment measurements did not exist at the time of the design of this experiment. Fortunately, any controller-structural instabilities would be bounded and non-destructive due to the limited dynamic pressure of the class B envelope, and the rate limits of the aileron actuators. Therefore, there were no safety concerns related to control-structural interactions. There were, however, mission success implications to undesired SMI responses. As such two configurations were certified for flight: one with no filtering on the strain feedbacks to explore any unmitigated SMI issues, and one with a simple first order low pass filter at 5 Hz to provide a backup configuration if the unfiltered configuration exhibited undesirable behavior. The sufficiency of this approach and the applicability of traditional SMI analysis techniques is addressed in Ref. 11.

An additional concern for these feedback parameters is time delay. This control allocation scheme is sensitive to time delay on the strain sensors. Intuitively, this sensitivity makes sense from an overload prevention standpoint. With fast control surfaces, significant delay in the strain feedback sensors will cause the control surface command to lead the measured load and result in the load overshooting the target, causing structural overload and large magnitude growing oscillations. The sensitivity to time delay is a function of the shape of the load constraint in the cost function and the maximum rate of the control surfaces. For slow surfaces and constraints that do not have steep slopes the allocation scheme is not very sensitive to time delay on the load feedback. For this experiment the ailerons were very fast control surfaces, and very steep cost functions were utilized, so time delay for some of the configurations was a significant concern. The time delay margin is easily verified in the simulation by progressively adding more time delay to the feedback signals and noting the point at which instability occurs for a given configuration.

VI. Simulation Results

A detailed study in the high-fidelity hardware-in-the-loop piloted simulation was completed as part of the development and certification process for the flight experiment. The simulation contains high-fidelity models of mass properties, aerodynamics, and actuators that have been verified with a significant amount of flight data. The simulation interfaces with real flight hardware; mission control computers, production and research flight control computers, and

research instrumentation processing hardware. One discrepancy between the simulation and the flight implementation is that in the simulation there is an additional 30 ms of delay on the feedback of the strain parameters to the research control computer. This additional time delay is a result of the hardware implementation of the instrumentation system in the simulation. It caused a few performance issues in the simulation that were not present in the flight implementation; however, the simulation still provided valuable insight and verified the soundness of the overall approach prior to flight.

A. Results for the Best Cost Function Configuration

The hardware-in-the-loop nonlinear simulation was used to characterize each of the cost function tuning parameters. Based on these piloted evaluations the configuration that had the most desirable characteristics was $n=20$, and $\epsilon=0.0001$. Figure 12 shows the performance of the algorithm during a 360° roll maneuver with this configuration and no filtering on the strain feedback parameters. The maximum roll rate achieved for both limiting on and limiting off with similar roll stick commands is nearly equivalent. In Fig. 12 the left aileron exceeds 10,000 in-lb of hinge moment without the limiting functionality engaged, but that the allocator enforces the 10,000 in-lb design limit with the load constraint engaged.

There are pronounced oscillations in the left hinge moment with limiting engaged in Fig. 12. These oscillations are the result of the 30 ms of simulation delay on the strain feedback parameters. The oscillations disappear when the strain models are implemented in the research flight-control computer, thus bypassing the simulation instrumentation system and eliminating the associated delay. The shape of the load constraint also affects this oscillatory behavior. Steeper load constraints (larger n values) increase the loop gain near the load limit. The higher the loop gain at a particular load value combined with sufficiently high time or phase delay exacerbates this undesired oscillatory behavior.

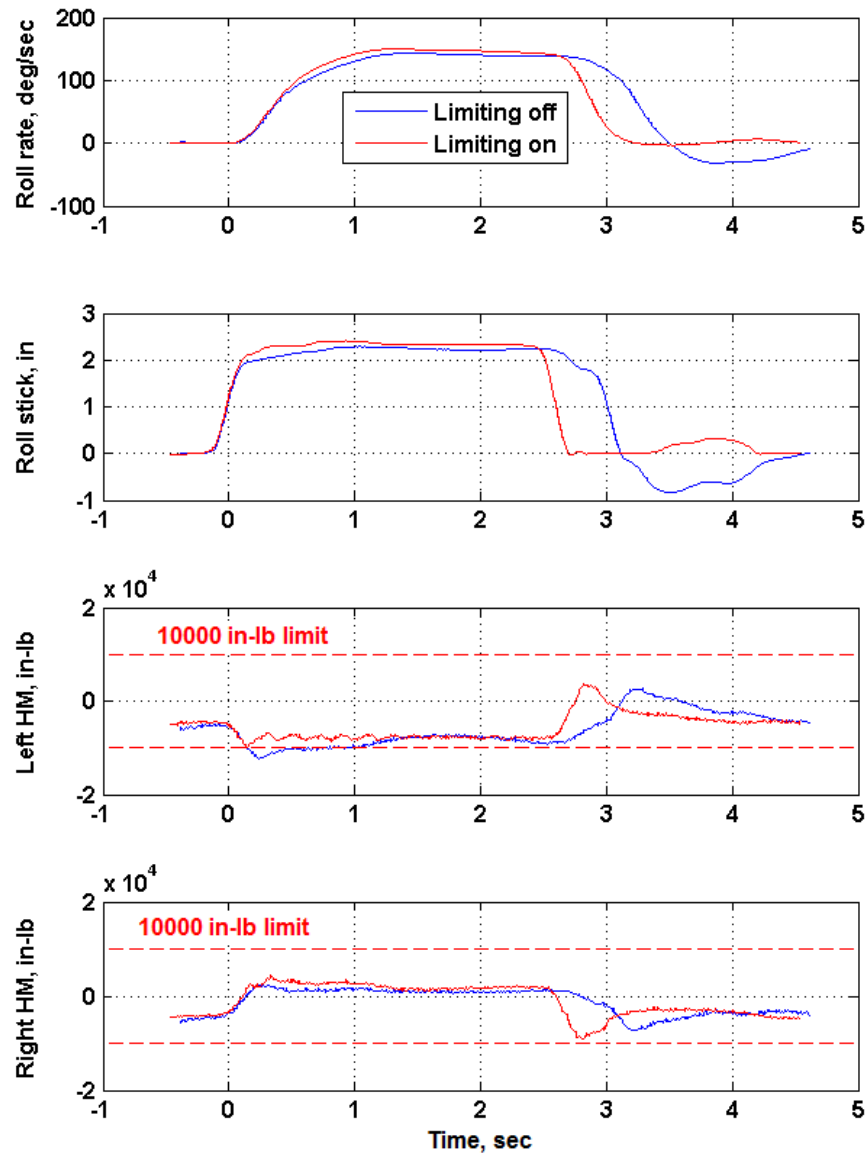


Figure 12. Comparison for limiting off versus on for a 360° roll.

Figure 13 shows the control surface usage for the same 360° roll maneuver. The peak differential aileron usage is reduced for the load limiting on case in accordance with the hinge-moment limit, and the other differential surface commands increase to compensate. Producing similar roll performance with less aileron available places an increased demand on the differential stabilator and the differential trailing-edge flaps.




Figure 13. Control surface command comparison for a 360° roll limiting off versus limiting on.

Figure 14 compares the converged value of the cost function and the number of iterations required to converge for the same 360° roll with both limiting on and off, and Fig. 15 shows each of the individual contributions to the converged cost function. Figure 15 shows that the dominant cost is the trim constraint. The very low values of the tracking cost suggest that good tracking is achieved for all cases. It can be seen that load limiting generally tends to increase the converged value of the cost function, and that load limiting on the allocator has to allow more tracking error and trim cost in order to minimize the cost associated high aileron loads near the specified limit. It can also be seen that more iterations are required for convergence of the algorithm when the load constraint plays a significant role in the total cost function.




Figure 14. Cost function comparison for 360° roll limiting off versus limiting on.




Figure 15. Individual cost function component comparison for 360° roll limiting off versus limiting on.

Figure 16 shows the individual values of the cost function contributions from one time step from the same 360° roll maneuver. It can be seen that for iterations 1 and 2 the dominant contribution to the cost function is the load constraint. As the iterations progress, the load cost value decreases, and the algorithm trades tracking error and trim cost in order to minimize load. Overall, after each iteration the value of the cost function is reduced. The norm of the vector of first derivatives of the cost function with respect to the surface commands can also be seen in Fig. 16. Each iteration drives the norm of the first derivative closer to zero as desired.




Figure 16. Cost function behavior at maximum iteration count for 360° roll limiting on.

Figure 17 shows the performance of the same configuration during a 2.5-g level turn. As the angle of attack increases so do both the average value and the noise level on the aileron hinge moments. Without the limiting engaged, the peak value of the left hinge moment exceeded the specified limit of 10,000 in-lb. With limiting activated, the control law is able to prevent the measured values from exceeding the desired limit. However, there is a significant amount of surface motion commanded as a direct result of the simulated surface buffet at higher angles of attack. This buffet represents a real load phenomena, as such simply filtering it out is not desirable for this application. The current system is unable to actively smooth the buffet phenomena due to insufficient bandwidth and excessive time delay. Additionally, for an operational system the amount of dither on the control surfaces would most likely be unacceptable. Smooth surface command responses need to be balanced with the responsiveness of the load limiting algorithm.




Figure 17. Comparison for limiting off versus on for a 2.5g level turn.

Figure 18 shows the converged value of cost function and the number of iterations required to reach convergence for the 2.5-g level turn. With limiting on, the largest values of the cost function are driven primarily by the load constraint. These large values are the direct result of the simulated buffet at high angle of attack. Since this buffet represents real aileron hinge moment, it is desirable for the limiting to respond by commanding the surfaces to reduce the load contribution to cost function. Figure 19 provides a little more insight into the individual contributions to the overall cost function for both limiting on and limiting off. Figure 20 illustrates the behavior of the iterative scheme when the measured load is high due to the buffet. The initial value of the load constraint cost function is very large. As the algorithm iterates, the value of the trim cost constraint increases, but only slightly and the load cost constraint is reduced several orders of magnitude while keeping the tracking error constraint near zero.

Comparing the cost function values from the 360° roll in Fig. 14 to those for the level turn in Fig. 18, it can be seen that the cost function values for the level turn are an order of magnitude less than the roll maneuver. However, many more iterations are required to reduce the load constraint enough to get these lower values of the cost function. For both roll and elevated g maneuvers, high values of the load constraint cause the algorithm to require more iterations to converge to an acceptable degree.




Figure 18. Cost function comparison for a 2.5-g level turn limiting off versus limiting on.




Figure 19. Individual cost function component comparison for a 2.5-g level turn limiting off versus limiting on.




Figure 20. Cost function behavior at maximum integration count for a 2.5-g level turn roll limiting on.

B. Effect of Time Delay on Load Feedback Measurements

Robustness to time delay is an important aspect of any feedback control application. The load measurements for this experiment were known to have 0.0063 seconds of transport delay on the aircraft; however, due to limitations in the simulation hardware, the transport delay in the simulation was estimated to be 0.03 seconds. Additional time delay on the critical load feedbacks results in poorer load limiting performance due to the fast response of the control surfaces. Additional time delay causes the measured load to overshoot the specified limit. The overshoots result in increased oscillatory behavior as well. The overshoots can be prevented by tuning the cost function to allow for more margin before the load constraint is activated (tune n and γ); however, activating the load constraint at lower loads reduces the control surface deflections that can be allocated to the command tracking objective and adversely affects vehicle performance. Figure 21 compares the simulation results with the nominal simulation time delay of 30 ms (blue) to a simulation run with an additional 32 ms of time delay (red). It illustrates the overshoot tendency and the increased magnitude of the load oscillation, and shows a reduction in the frequency of that oscillation. All of those

effects are consistent with a reduction in the phase crossover frequency and a reduction in phase margin with increasing time delay on load feedback parameters.




Figure 21. Comparison of controller performance with 32 ms of additional time delay on the strain feedback.

C. Effect of Load Constraint on Roll Performance and Convergence

The load constraint tuning affects the control surface usage and how close to the 100% load limit the algorithm will allow the load measurements to get without sacrificing command tracking performance to alleviate the measured aileron loads. As was discussed in section IV, steeper constraints allow the measured load to get very close to the specified limit before sacrificing tracking performance. Figure 22 shows how the lower order equivalent system (LOES) roll mode performance varies with both load constraint steepness and specified aileron hinge-moment limit. It can be seen that for decreasing load limits the roll mode time constant and the roll rate per stick command values both indicate a reduction in roll performance. This behavior is expected as lower hinge moments allowable translate directly to lower achievable roll moments. The roll performance drop-off is less severe for the steepest load constraint. It should be noted that the frequency response of the OCLA controller and LOES is a linear analysis (see Ref. 12

Hodgkinson for a description of LOES applicability and analysis techniques). Figures 23 and 24 show time domain comparisons of $n=20$ and $n=4$ for a 360° roll maneuver. It can be seen from these figures that $n=20$ allows the algorithm to get closer to the full aileron command available, utilizing less differential stabilator and trailing-edge flap, and $n=20$ achieves a faster roll rate onset and higher maximum roll rate. The oscillatory behavior for $n=20$ in Fig. 23 is the result of the time delay on the strain signals.

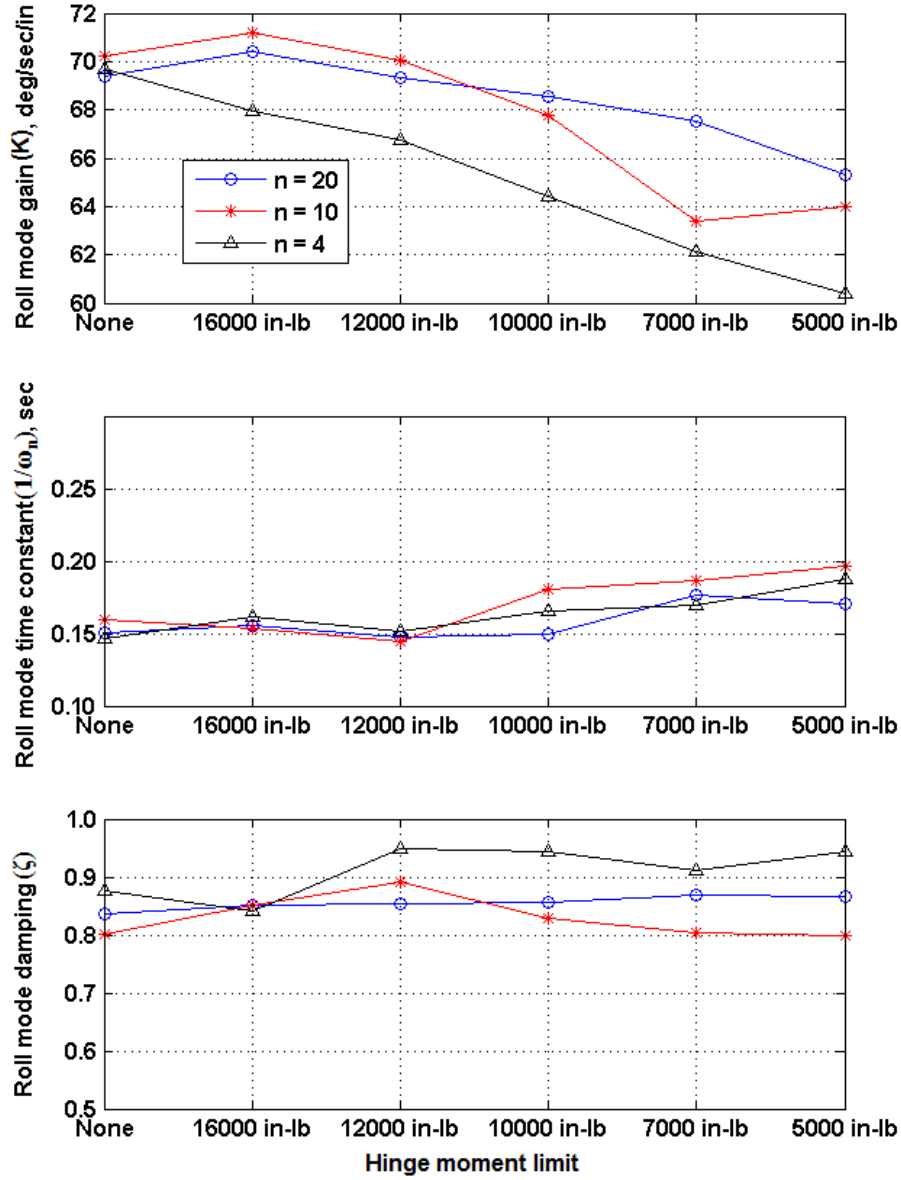


Figure 22. Roll mode parameters $\frac{p}{dap} = \frac{K\omega_n^2 e^{-0.08s}}{s^2 + 2\zeta\omega_n s + \omega_n^2}$ ($\epsilon = 0.0001$).




Figure 23. 360° roll comparison $n=20$ versus $n=4$.




Figure 24. Surface usage comparison for 360° roll n=20 versus n=4.

Figure 25 compares the converged value for the cost function and the number of iterations required to reach convergence for the 360° roll (Fig. 23). n=4 results in higher values of the cost function, but it does not require as many iterations to converge. The higher iteration count for n=20 is attributed to the fact that at higher loads the cost function is highly non-linear for n=20, which complicates convergence for a linear Newton-Rapson optimization routine. Figure 26 shows the individual contributions to the cost function. For n=4, tracking performance is sacrificed and the trim cost is higher because the load constraint is activated at lower measured values of the hinge moments. This early activation results in less aileron command and higher stabilator and trailing-edge flap deflections.




Figure 25. Cost function value and iteration count comparison for 360° roll $n=20$ versus $n=4$.

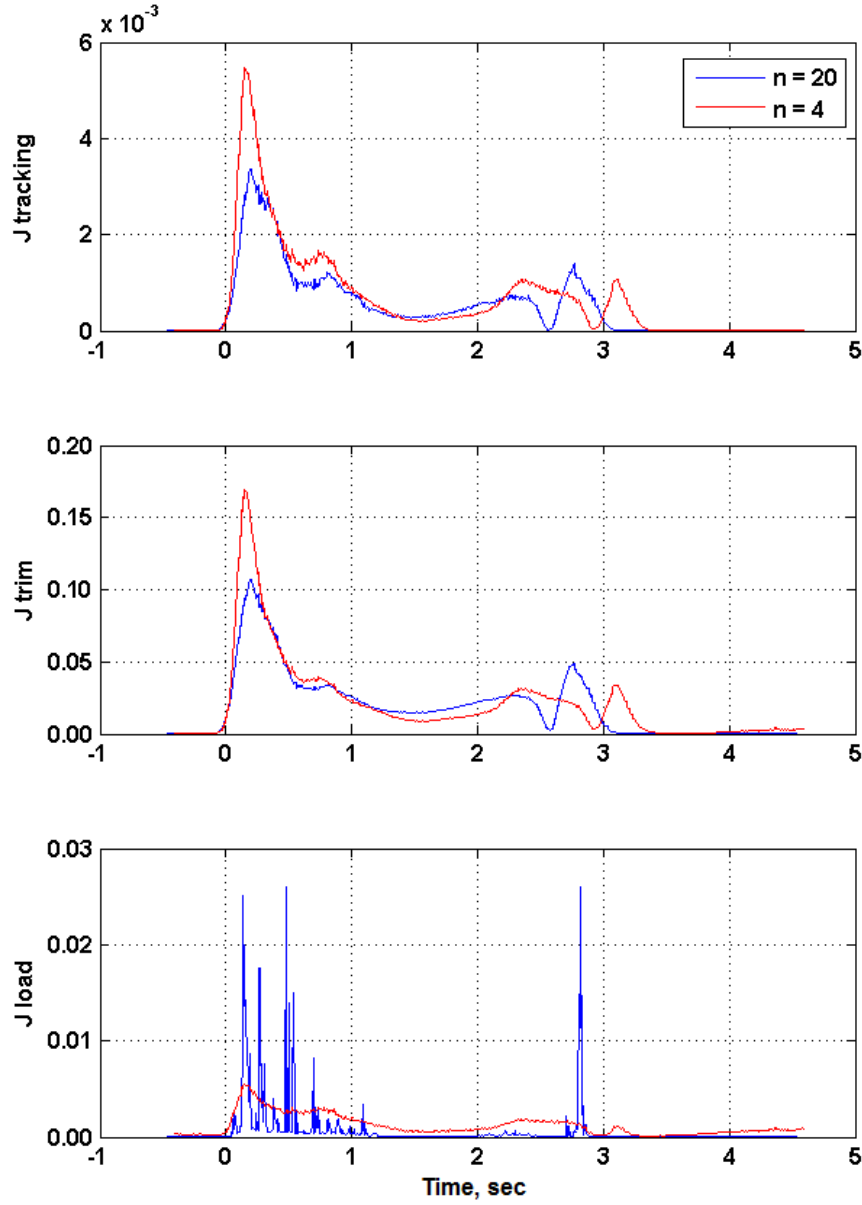


Figure 26. Individual cost function contributions comparison for 360° roll $n=20$ versus $n=4$.

D. Effect of Trim Constraint Weight on Roll Performance

The trim constraint is used to attract the surfaces to a predefined set of trim positions. This constraint effectively resists moving the surfaces away from their trim positions. Intuitively it makes sense that if this constraint is weighted strongly then the algorithm will sacrifice tracking performance in order to maintain the surface trim. As discussed in section IV keeping the weight, ϵ , as small as practical is desirable for achieving good command tracking and handling qualities. Figure 27 shows that as the trim constraint weight is increased the roll mode properties are adversely affected. Large values of ϵ result in lower roll rate per stick command gains, and an over damped roll response with a mostly constant roll mode time constant. These adverse effects related to large values of the trim constraint weight can also be seen in Fig. 28. The same maximum roll rate is achieved for both values of ϵ ; however, the larger value exhibits a larger rise time. Figure 29 compares the individual contributions to the cost function for the 360° roll

maneuver shown in Fig. 28. Larger ε results in larger values of the trim cost, which results in larger values of the tracking error cost. The larger values of the trim weight result in degraded tracking performance.

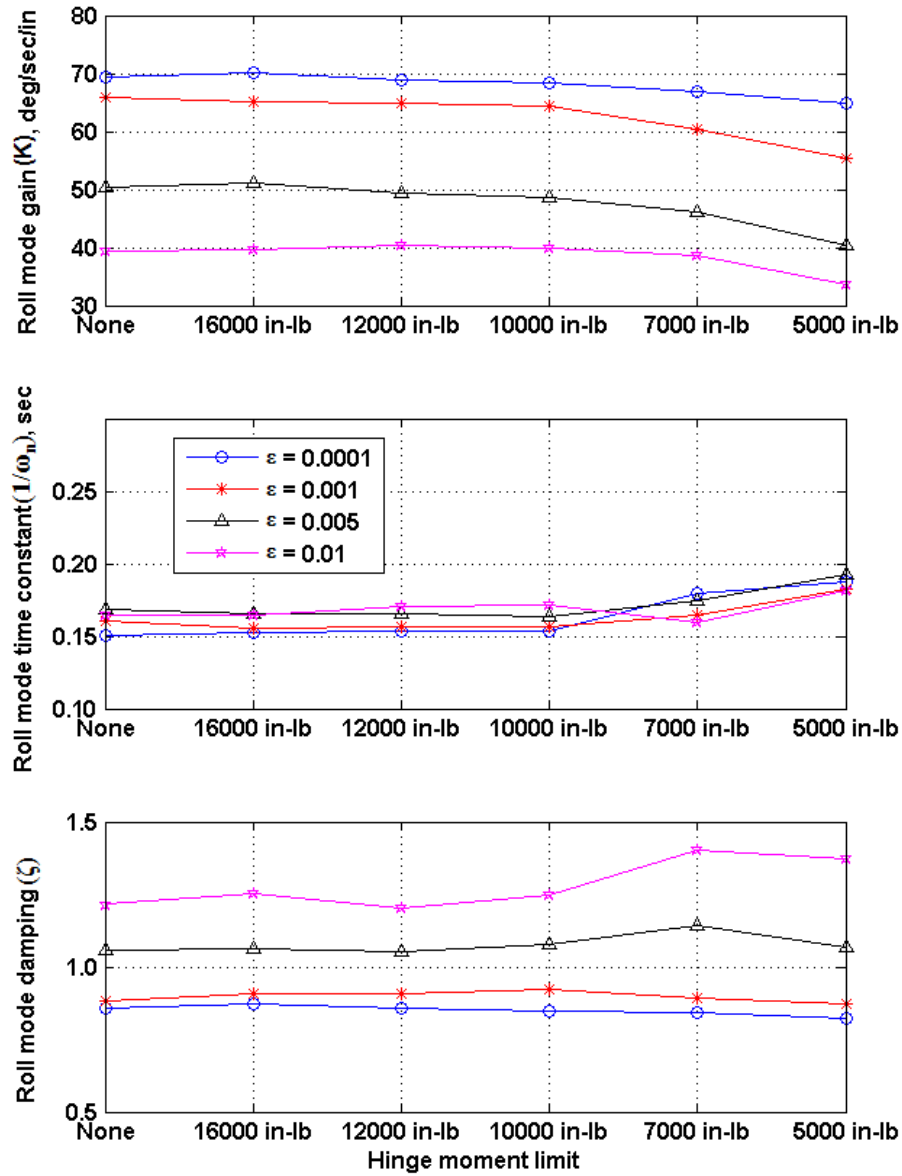


Figure 27. Roll mode $\frac{p}{dap} = \frac{K\omega_n^2 e^{-0.08s}}{s^2 + 2\zeta\omega_n s + \omega_n^2}$ parameters ($n = 20$).

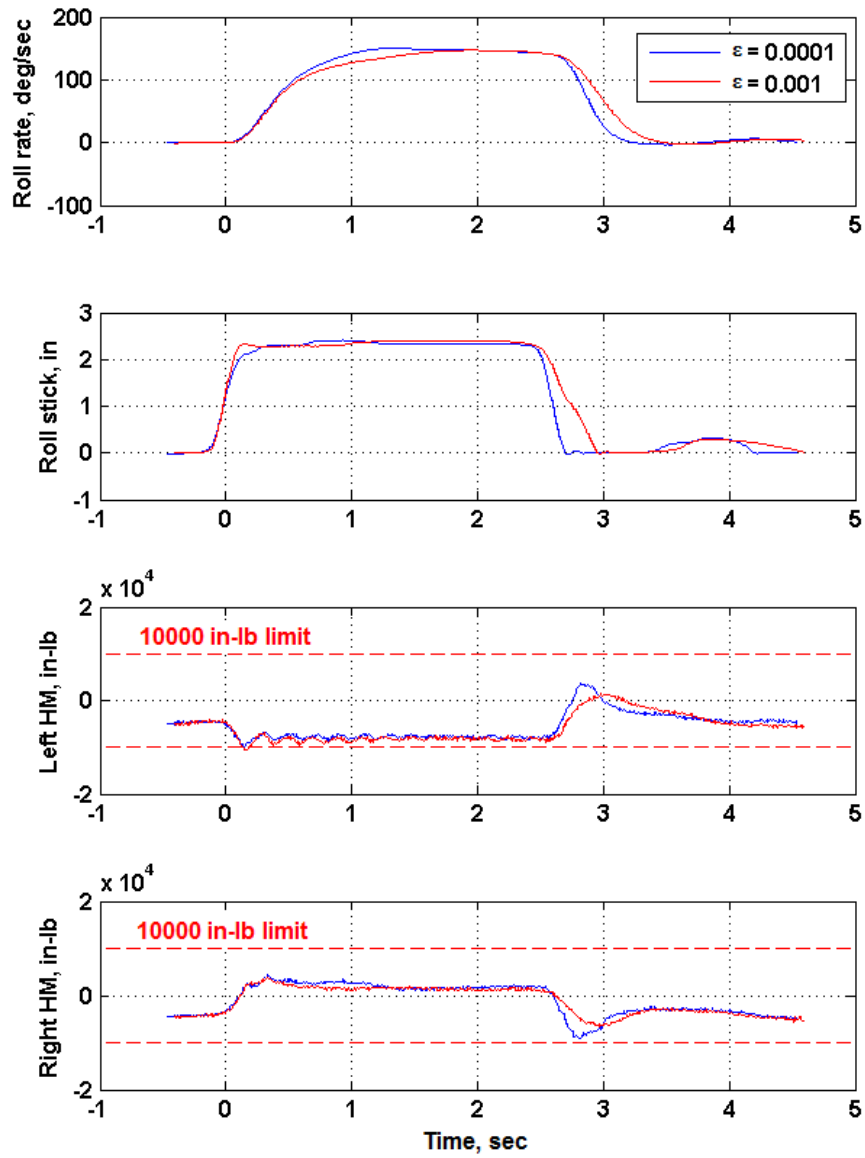


Figure 28. 360° roll comparison $\epsilon=0.0001$ versus $\epsilon=0.001$.




Figure 29. Individual cost function contributions comparison for 360° roll $\epsilon=0.0001$ versus $\epsilon=0.001$.

E. Flight Condition Variation for Roll Performance

Achieving good performance across a range of flight conditions is an important aspect of any control system. The dynamic inversion architecture with constant command reference models as implemented for this experiment should show minimal variation in roll response with changing flight conditions. Figure 30 compares LOES roll mode properties at FC6 and FC11. The roll mode is very similar for load limits which permit free use of the ailerons; however, as the aileron hinge-moment limit becomes more restrictive it can be seen that roll mode gain per stick command drops off more steeply for the slower flight condition. Intuitively, this drop in roll mode gain makes sense because at lower dynamic pressures larger control deflections are necessary to produce equivalent roll moments. As the ailerons are removed from the command tracking objective there just is not enough control power to generate the commanded roll moments.

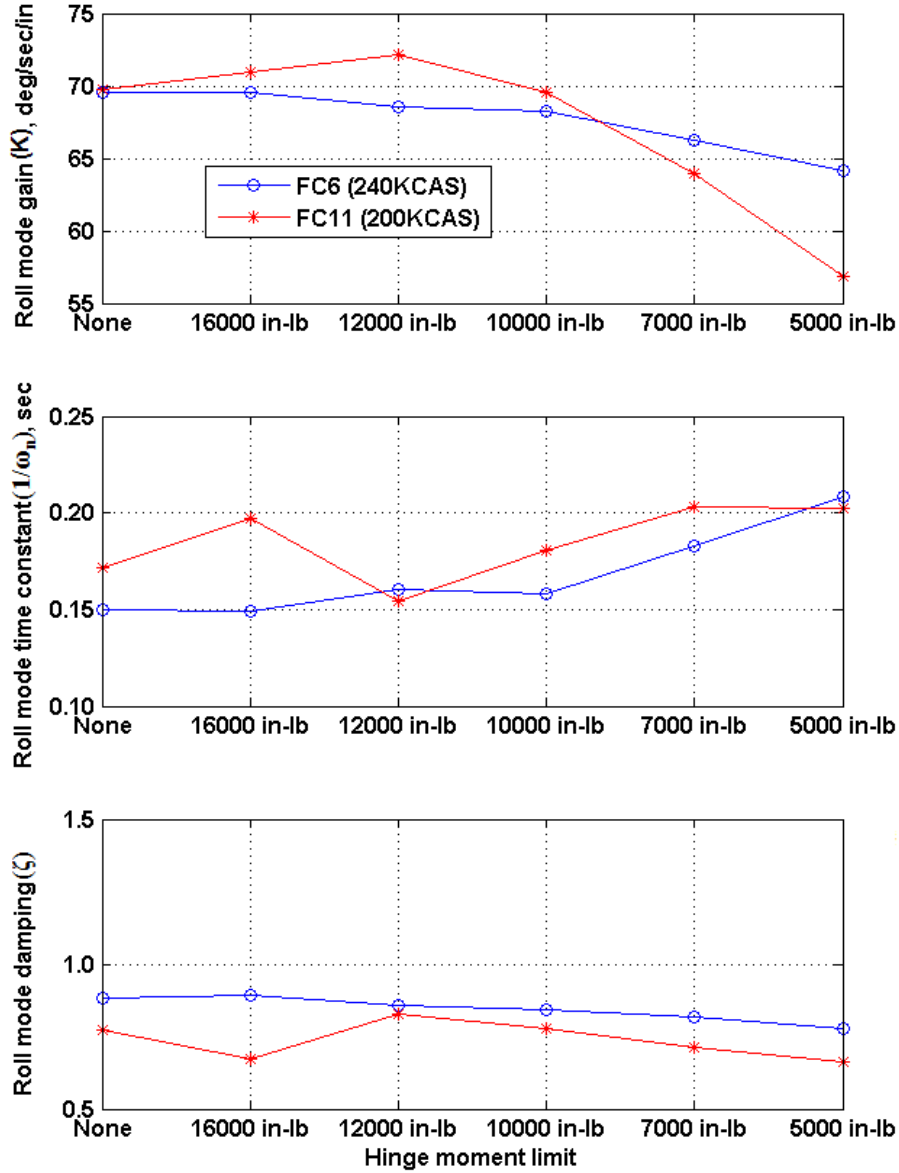


Figure 30. Roll mode $\frac{p}{dap} = \frac{K\omega_n^2 e^{-0.08s}}{s^2 + 2\zeta\omega_n s + \omega_n^2}$ parameter comparison for flight condition ($n = 20$, $\varepsilon = 0.0001$).

VII. Conclusions

The goal of this research was to explore a practical approach to limiting critical aircraft loads while maintaining control performance. The approach described in this paper accomplished this goal by developing a cost function and optimization approach for an optimal control allocation. The cost function contains elements for command tracking, specifying vehicle trim, and critical load limiting. The scheme provided hard load constraints at high loads, but allowed the free use of all control surfaces at low sensed loads. The formulation allows for the implementation of a functional hard load constraint on the aircraft loads without requiring a unique mapping between loads and surface commands. The allocation scheme redistributes control commands and loads away from a structure that is near limit loads, and to control surfaces with remaining control capability and structural margin. The approach was shown to limit the loads to acceptable levels 100% of the time for all combinations of the cost function tuning parameters. The load limiting

was shown to have minimal effects on the aircraft control performance when well-tuned. The control technique displayed the ability to trade tracking error, and aircraft trim penalties for load suppression even with very restrictive load limits.

Tuning of this cost function was found to be straight forward and intuitive. The trim constraint provides the designer with the flexibility to specify where in the overdetermined space they want the aircraft surfaces to trim (u_p). The H matrix lets the designer set how aggressively each surface can be used for the various control objectives. Keeping the size of the trim constraint small relative to the other constraints is important for achieving desirable performance, and can be accomplished via the parameter ϵ . The load constraint steepness (n) made it possible to aggressively limit loads when they are high while keeping the load constraint functionally dormant at low load. Tuning of this parameter in conjunction with the overall magnitude of the load constraint (γ) was found to be important to get desirable load limiting and command tracking performance. Finally, a linear assumption for control surface commands to incremental loads within the load constraint was found to be sufficiently accurate.

The cost function developed represents a continuous and differentiable convex space which means that it has a global minimum. The fact that the cost function forms a convex space facilitated the use of a simple Newton-Rapson optimization approach which made convergence predictable and efficient. The addition of a relaxation factor to the optimization technique improved the convergence properties further and resulted in the cost function value monotonically decreasing with each iteration.

The load sensors chosen for the experiment with the redundancy management system implemented proved to be sufficient for the purposes of this experiment. The mostly linear models for these sensors developed from flight-test data worked exceptionally well for the simulation characterization of the control laws.

The experiment did uncover a number of issues that bear further research. The first is a need for robust redundant instrumentation of both critical loads and control surface position with high sample rate and minimal delay. The control law configured to aggressively limit loads is very sensitive to time delay in both control surface measurements and the load measurements. Steep load constraints were found to be the most susceptible to performance issues with high time delays and also increased the time required for the optimization to converge. Aerodynamic buffet was found to be especially challenging to account for, when attempting to suppress critical loads. Increases in controller bandwidth and reductions in time delay may help address this challenge. Finally, additional work needs to be done to prove out optimal control allocation with load limiting before it can be utilized for an application without a robust backup control algorithm to fall back on. However, even with these challenges the optimal control allocator as outlined in this paper has been shown to be promising for limiting critical aircraft loads while maintaining control performance.

References

- ¹Frost, S. A., et al., "Application of Structural Load Feedback in Flight Control," AIAA-2011-6288, 2011.
- ²Bodson, M., and Frost, S. A., "Load Balancing Control Allocation," *Journal of Guidance, Control, and Dynamics*, Vol. 34 No. 2, March-April 2011.
- ³Bodson, M., "Evaluation of Optimization Methods for Control Allocation," *Journal of Guidance, Control, and Dynamics*, Vol. 25, No. 4, July-August 2002.
- ⁴National Transportation Safety Board Aircraft Accident Report, "In-Flight Separation of Vertical Stabilizer American Airlines Flight 587 Airbus Industrie A300-605R, N14053 Belle Harbor, New York November 12, 2001," NTSB/AAR-04/04 PB2004-910404 Notation 7439B, 2004.
- ⁵Cooper, G. E., and Harper, R. P. Jr., "The use of Pilot Rating in the Evaluation of Aircraft Handling Qualities," NASA TN D-5153, 1969.
- ⁶Pavlock, K., "Full-Scale Advanced Systems Testbed: Ensuring Success of Adaptive Control Research Through Project Lifecycle Risk Mitigation," *Society of Flight Test Engineers*, SFTE 2011 42-12, 2011.
- ⁷Miller, C. "Nonlinear Dynamic Inversion Baseline Control Law: Architecture and Performance Predictions," AIAA 2011-6467, 2011.
- ⁸Miller, C. "Nonlinear Dynamic Inversion Baseline Control Law: Nonlinear Dynamic Inversion Baseline Control Law: Flight-Test Results for the Full-scale Advanced Systems Testbed F/A-18 Airplane," AIAA 2011-6468, 2011.
- ⁹Schaefer, J., and Brown, N. A., "Peak-Seeking Optimization of Trim for Reduced Fuel Consumption: Architecture and Performance Predictions," AIAA 2013-4856, 2013.
- ¹⁰Brown, N., and Schaefer, J., "Peak-Seeking Optimization of Trim for Reduced Fuel Consumption: Flight-Test Results," AIAA 2013-5171, 2013.
- ¹¹Miller, C. J., and Goodrick, D., "Optimal Control Allocation with Load Sensor Feedback for Active Load Suppression, Flight Test Performance," AIAA 2016-#### (to be published), 2017.
- ¹²Hodgkinson, J., and LaManna, W. J., "Equivalent System Approaches to Handling Qualities Analysis and Design Problems of Augmented Aircraft," AIAA-1977-1122, 1977.